

Techniques for Collision Prevention, Impact Stability, and Force Control by Space Manipulators

Richard Volpe

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109

1 Introduction

The field of space robotics can be readily divide to planetary and zero-gravity operations. While the harsh environments of other planets will surely require robust robotic hardware, the algorithms controlling this hardware are not likely to be different in kind from earth-based controllers. Therefore, it is usually the arena of zero-gravity operations where special control algorithms are developed for space robotics [54, 43, 73, 74, 71]. Among the pertinent issues that this research addresses are: six-degrees-of-freedom mobility, zero friction motion, energy (thrust) minimization, large inertias, flexible structures, etc. Many of these have terrestrial analogs, especially in the field of underwater robotics (as is demonstrated by the utility of buoyancy tanks for astronaut mission training). Particularly, robot manipulation in space has a large overlap with its terrestrial counterpart.

Within this overlap of zero-gravity and terrestrial robotics, there are three main issues: unconstrained motion, stability during the contact transition, and force controlled manipulation of the environment. If a space robot is unattached to its environment, the first two of these research areas map closely to the problems of mobile ground robots and underwater vehicles. In this case, the main problems are path planning, obstacle avoidance, and rendezvous and docking. Force control is not pertinent, because any forces exerted between the robot and its environment will tend to repel each away from the other. While this is especially true in space, it can also be a practical matter for mobile robots on land and in water. Therefore, if force control is to be applied, the robot should attach itself to the environment, making a continuous kinematic chain. (Constant force could also be applied by thrusters, wheels, or propellers, but it is inefficient and will cause a net acceleration if the environment is not grounded.) The attachment of the robot to the environment is typically achieved through slow docking followed by base attachment, or grasping by one arm of a multi-arm system.

Once the robot and its environment are coupled, the manipulation control issues are essentially the same for space, ground, or water robots. In this case all three problems of robot collision free motion, contact transition, and force control are important. Collision free motion is often more difficult with the constraint of the attached base [58]. The attachment does, however, allow the control of interaction forces, in the form of impact control and accurate force trajectory following on the contacted surfaces.

This chapter discusses several control strategies which have successfully addresses these problems of real-time collision free motion through the environment, reduced velocity approach of surfaces to be contacted or docked with, impact control, and force control. Each of these techniques will be reviewed, and analysis, simulation, and experimentation will be presented. Finally, there will be a complete discussion of the implementational issues for all of these strategies.

Section 2 discusses the use of artificial forces provided by superquadric artificial potential functions. These functions can assume the shape of a wide variety of objects and provide complete repulsion from objects, or just enough repulsion to slow an approach for safe impact. To address the control of the robot after this impact, Section 3 reviews a method of impact control that provides stable, bounceless contact. A full understanding of this control strategy requires a presentation of the system model, which is provided in Section 4. Further, the impact controller analysis and experimentation indicates that two seeming disparate control schemes, second order impedance control and proportion gain explicit force control, are essentially equivalent. This issue is discussed in detail in Section 5, and followed by a discussion of other explicit force control schemes in Section 6. Finally, Section 7 provides a detailed discussion of the implementational consideration needed to understand the behavior of all of these control schemes, and to make them work in practice.

2 Artificial Forces

Moving amid the space station is essentially the same problem as moving through planetary terrain, or moving through a factory workcell. All require the avoidance of obstacles to reach the goal location. Research in obstacle avoidance can be broadly divided into two classes of methods: *global* and *local*. Global methods rely on the description of the obstacles in the configuration space of a manipulator [53, 33, 47]. Local methods rely on the description of the obstacles and the manipulator in the Cartesian workspace [27, 4, 31].

Global methods require that two main problems be addressed. First, the obstacles must be mapped into the configuration space of the manipulator [33]. Second, a path through the configuration space must be found for the point representing the manipulator. Two techniques are used to generate these paths: *geometric searches* and *artificial forces*. The geometric search technique relies on an exhaustive search of the unoccupied configuration space for a continuous path from the start point to the goal point [53, 35, 34, 50, 16]. If a path exists, it will be found. If multiple paths are found, the best may be chosen. The artificial force technique surrounds the configuration space obstacles with repulsive potential energy functions, and places the goal point at an global energy minimum [42, 47, 38, 67]. The point in configuration space representing the manipulator is acted upon by a force equal to the negative gradient of this potential field, and driven away from obstacles and to the minimum.

Global methods have several disadvantages. The algorithms necessary for global methods are computationally intensive. Also, the computational costs increase quickly as a function of the manipulator's degrees-of-freedom: at least exponentially for geometric search techniques, and at least quadratically for artificial force techniques [47]. Thus, they are suited only for off-line path planning and cannot be used for real-time collision avoidance. An immediate consequence is that global algorithms are difficult to use for collision avoidance in dynamic environments, where the obstacles are moving in time. Also, using global algorithms it is very difficult to describe complicated motion planning tasks such as those arising when two manipulators cooperate.

A viable alternative to global methods is provide by local ones [27, 4, 31, 57]. Local methods also employ the use of artificial forces like those discussed previously. However, unlike configuration space forces, local forces are expressed in the Cartesian workspace of the manipulator. Collisions with objects are prevented by surrounding them with repulsive potential functions, and the goal point is surrounded by an attractive well. These potentials are added to form a composite potential which imparts forces on a model of the manipulator in Cartesian space. Torques equivalent to these forces cause the motion of the real manipulator.

The main advantage of local techniques is that they are less computationally demanding than

global ones, permitting their use in real-time control. Further, they provide the necessary framework to deal with changing environments and real-time collision avoidance. When used with a teleoperated manipulator, local artificial forces also provide low level collision avoidance, while high-level path planning of the manipulator is performed by the human operator.

2.1 Model-Based Potential Energy Functions

The major interest in artificial force models has been in realizing obstacle avoidance schemes [31, 30, 39, 19, 58]. These schemes require the addition of attractive and repulsive potential energy functions. An attractive potential well is generally a bowl shaped energy well which drives the manipulator to its center if the environment is unobstructed. However, in an obstructed environment, repulsive potential energy hills are added to the attractive potential well at the locations of the obstacles, as in Figure 1. The addition of attractive and repulsive potentials provides the artificial forces which enable obstacle avoidance.

The assignment of potential energy values to the isopotential surfaces determines the repulsive nature of the function. Two possibilities exist: avoidance and approach functions. The avoidance function has a potential energy value at the surface of the object which is larger than the initial kinetic energy of the manipulator. Thus, an energy barrier is established which cannot be surmounted. The easiest way to ensure that the potential energy barrier is large enough is to set the potential function to infinity at the object surface.

We have previously proposed a second type of artificial potential energy function for approach [58]. Instead of a potential of infinity at the object surface, the function goes smoothly to a finite value less than the kinetic energy of the manipulator. As the manipulator moves toward the object, it gains potential energy, loses kinetic energy, and slows down. Thus the approach potential provides deceleration forces that ensure a safe contact velocity at the surface. Once stable contact has been established, force control of the manipulator may begin.

Many proposed repulsive potentials have spherical symmetry [4, 30, 58]. These potentials are useful for surrounding objects with spherical symmetry, as well as singularities in the workspace. Also, when added to a spherically symmetric attractive well they will not create a local minimum [58]. But a spherically symmetric repulsive potential does not follow the contour of polyhedral objects. For instance, an oblong object surrounded by a sphere effectively eliminates much more volume from the workspace than is necessary or desirable. Potentials that follow the object shape were proposed to address the insufficiency of radially symmetric potentials [31, 27, 58]. A review of one of these schemes will be provide next.

2.2 Superquadric Isopotential Contours

We have proposed an artificial potential scheme based on the superquadric, a mathematical function which is employed in computer vision and object modelling techniques [5, 6]. This scheme provides obstacle avoidance capability for manipulators in an environment of stationary or moving objects, preventing end effector and link collisions with these objects. This local avoidance scheme provides obstacle avoidance capability without creating local minima.

The superquadric is a deformable parametric surface and is used in this scheme as the isopotential surface for the potential function. Since it is deformable, isopotential surfaces near the object may closely model the object, while surfaces further away can be spherical. These spherical surfaces prevent the formation of local minima when this function is added to a larger spherical attractive potential well.

To obtain isopotential contours that follow the object shape near the surface an object is surrounded with a superquadric [5, 6]:

$$K = \left[\left(\frac{x}{a} \right)^{2n} + \left(\frac{b}{a} \right)^2 \left(\frac{y}{b} \right)^{2n} \right]^{\frac{1}{2n}} - 1. \quad (1)$$

Figure 2 shows a plot of K at regular intervals with n varying from a very large value to a value near unity.

Since the parameter n must vary from infinity to one while K varies from zero to infinity, n is defined as:

$$n = \frac{1}{1 - e^{-\alpha K}} \quad (2)$$

where α is an adjustable parameter. Other definitions of n are possible, but this form is useful because it is related to the magnitude of the potential (as will be shown).

2.3 Repulsive Artificial Force Functions

The artificial force experienced is dependent on the form of the potential energy function assigned to the isopotential contours defined previously. We have utilized two types of repulsive energy functions: the avoidance function, and the approach function.

The Avoidance Function The avoidance function surrounds an object and prevents a manipulator from touching the object. This is true, independent of the manipulator's kinetic energy, and is ensured by setting the magnitude of the potential at the surface to infinity. Away from the surface, the energy values behave like natural potentials (e.g. electrostatic, gravitational, etc.) in their inverse dependence on distance. This is done with the Yukawa potential [9] which has K^{-1} dependence for short distance repulsion, but exponential decay at larger distances:

$$U(K) = A \frac{e^{-\alpha K}}{K} \quad (3)$$

Figure 3 shows this function with $\alpha = 1$ and $A = 1$ for a rectangle. The parameter α determines how rapidly the potential rises near the object and falls off away from the object. This rate is tied to the rate at which the ‘ n -ness’ of the ellipse changes as expressed in Equation (2). The parameter A acts as an overall scale factor for the potential. Large values of A will make the object have a spherical field of repulsive forces at large distances. Small values of A will allow the object to be approached much more closely. At this closer range, the isopotential contours will have large values of n and will approximate the shape of the object. Typically A is unity.

The Approach Function The approach function surrounds an object and decreases the approach speed of the manipulator as it moves toward the object. This is achieved by setting the value of the potential energy at the surface of the object to be slightly less than the initial kinetic energy of the manipulator. As the manipulator moves toward the object its kinetic energy is transformed to potential energy, and its velocity decreases. Setting the magnitude of the potential function at the surface less than the initial kinetic energy ensures that the manipulator will always reach the surface.

An appropriate approach potential has all of the attributes of the avoidance potential, but goes to a finite maximum value at the surface of the object. Therefore, far from the object, the form of

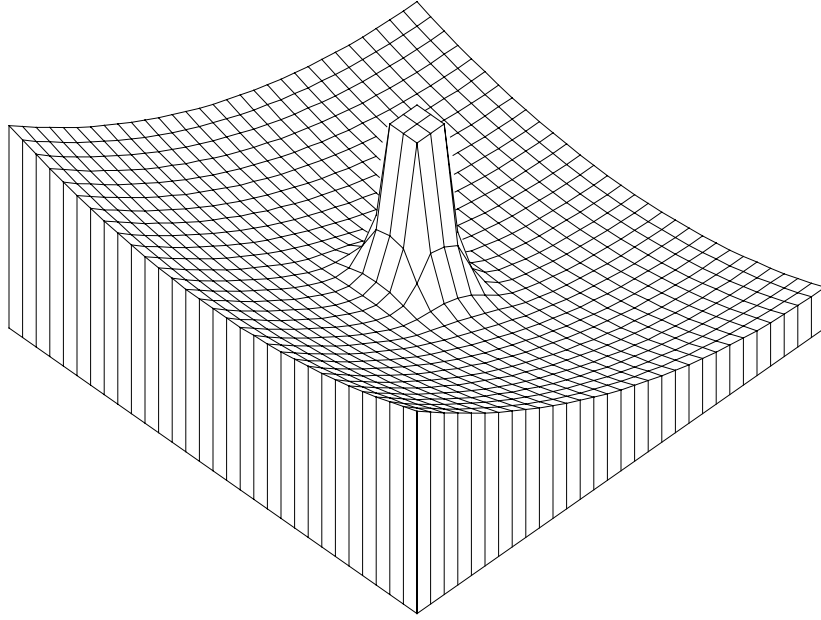


Figure 1: A repulsive potential added to an attractive well.

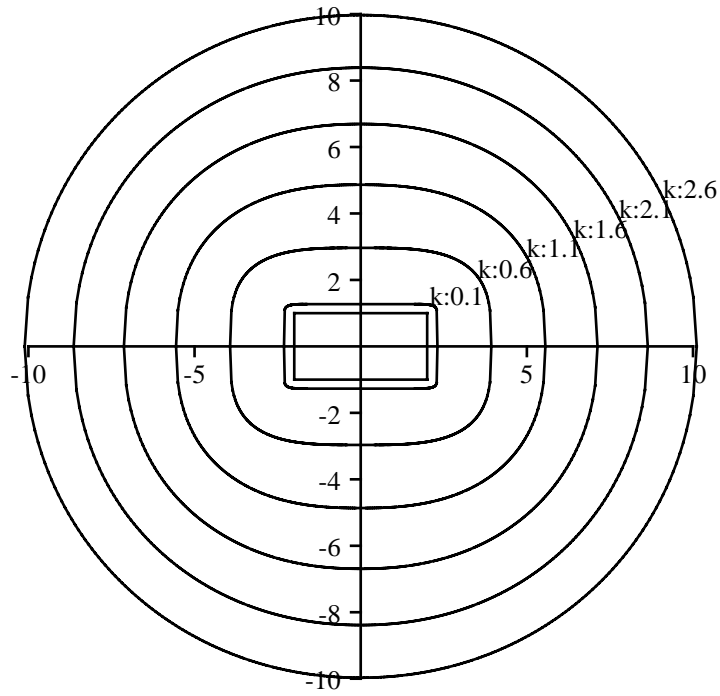


Figure 2: The isopotential contours for $K = 0.1$ to $K = 2.6$, and $\alpha = 1.5$.

the avoidance potential is used. However, closer to the surface the potential is Gaussian in shape, the slope smoothly changing to zero at the surface so that no artificial force is experienced when real contact with the environment is established. A general function form which remains valid for all values of α is:

$$U(K) = \begin{cases} \frac{A}{K} e^{-\alpha K}, & K \geq 1 \\ A \exp\left(-\alpha K^{1+\frac{1}{\alpha}}\right), & 1 > K \geq 0 \end{cases} \quad (4)$$

Figure 4 shows this function with $\alpha = 1$ for a rectangle.

2.4 Simulation and Experimentation

The collision prevention and approach capabilities provided by these functions are demonstrated by simulations of two and three link manipulators. Figures 5 shows simulation of successful avoidance of several objects by a three-link manipulator. Figure 6 show simulation of smooth approach of an obstacle by a two-link manipulator. Figure 7 shows experimental data of successful avoidance of several objects by the end-effector of a SCARA type manipulator.

These results rely on knowledge of the spatial relationship between the robot and the obstacle. For the test completed, this was calculated from a priori knowledge of the object position. However, the formulation directly supports measurement of the relationship by cameras or proximity sensors. This type of sensor based collision avoidance has been explored by several researchers [8, 66, 41, 40, 7, 2]. It remains an area of active research since the distance measurements tend to be sparse, noisy, inaccurate, and difficult to physically, electrically, and computationally integrate into the system. However, sensor-based knowledge of the environment promises to be the best way to avoid obstacles or prepare for intended impacts with it.

3 Impact Control

Even if the manipulator has be slowed by a repulsive approach force, switching from free-space motion to constrained force control has the significant problem of impact forces [44]. These forces can be very large, and can drive an otherwise stable controller into instability. While detrimental for terrestrial operations, in space large impact forces can have added severity. Large forces can repel the environment, or excite its modes of oscillation. Further, lightweight and delicate space hardware is more susceptible to damage, difficult to repair, and more costly to replace. Given these problems, robust impact control is extremely valuable.

Typically, it is the force control strategy that must deal with this transient phenomenon, since the large force does not occur until after contact has occurred. However, the natural elasticity of the impact, or the response of the force controller to the transient, can cause the manipulator to rebound from the environment. Thus, the manipulator is once again unconstrained. This phenomenon can establish oscillatory behavior or worse, drive the manipulator unstable. Obviously it is the goal of any controller to pass through this transitory period successfully, and have the manipulator stably exerting forces on the environment. The controller must, therefore, pass through the impact phase by attempting to maintain contact with the environment until all of the energy of impact has been absorbed. To maintain stability and contact during this phase, strategies for *impact control* will be reviewed in this section [65].

Previous research in force control has treated the impact phase as a transient that is dealt with by the same controller used to follow commanded force. The form of the force controller is typically an explicit force or impedance controller [70, 56]. In this section it will be shown that the best implementation of these strategies for force following is insufficient for impact control.

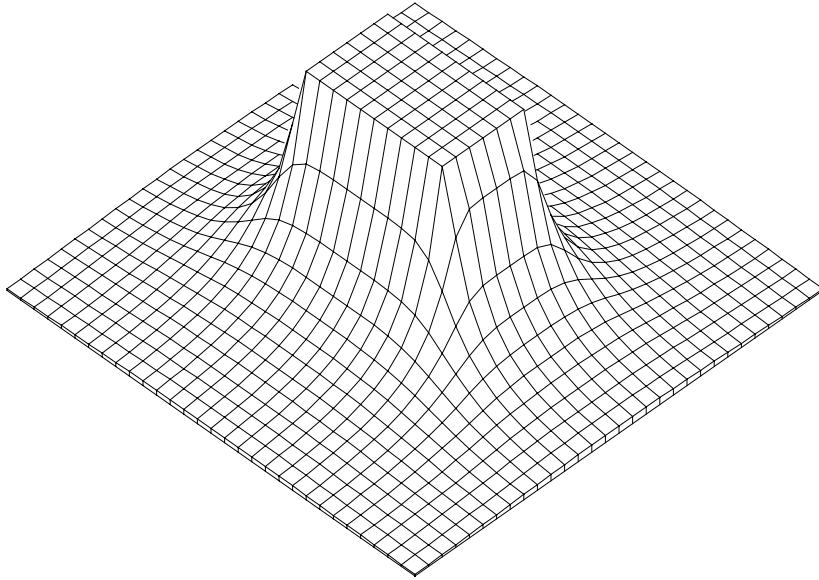


Figure 3: The avoidance potential for a rectangle with $\alpha = 1$ and $A = 1$. Large values have been truncated.

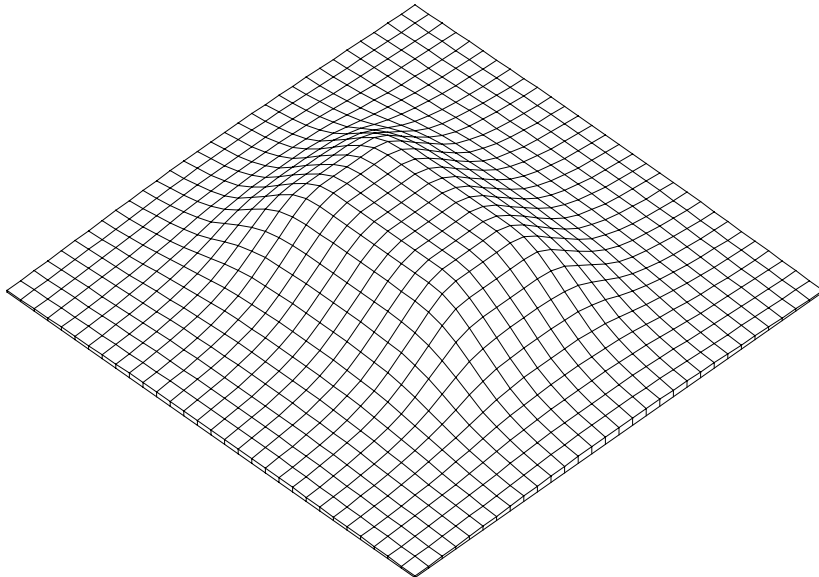


Figure 4: The approach potential function for a rectangle with $\alpha = 1$.

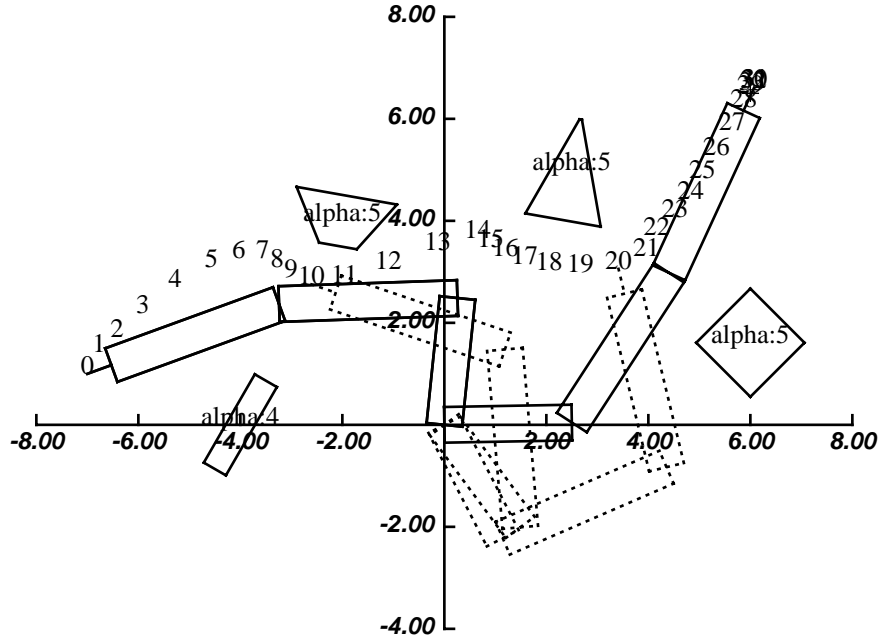


Figure 5: Successful navigation around four obstacles using superquadric avoidance potentials and a modified conical attractive well. The dotted manipulators are intermediate configurations.

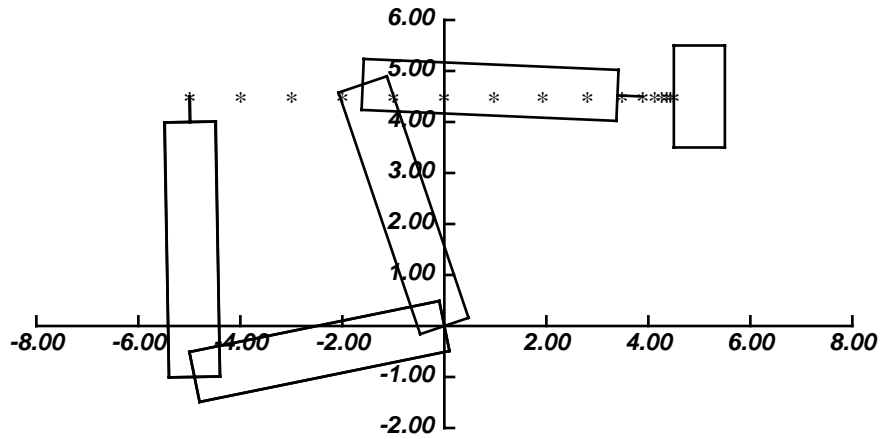


Figure 6: This figure shows successful approach and contact with a rectangle surrounded by the proposed approach potential. For this simulation there was no attractive point, but the end effector was position controlled in the y direction. The initial velocity was 1 unit/sec in the x direction. The contact velocity was 0.06 unit/sec.

But the impact controller presented here still fits into the same framework. To understand this, the previous schemes will be briefly discussed and their weaknesses revealed. Then our previously proposed impact control strategy will be presented in the context of explicit force control and impedance control. An analysis will explain how the strategy provides stability, and experimental results will demonstrate its effectiveness.

3.1 Previously Proposed Methods For Impact Control

Most previous work in impact control has not employed any changes in the force controller structure (variation of gains or controller type). Instead the impact phase is treated as a transient that must be dealt with by the force controller and the chosen gains, once contact has been established. Typically, modification of the control strategy has been attempted through active damping and/or passive compliance and damping.

Maximal Active Damping One proposed method of dealing with the impact problem is to employ maximal damping during the impact phase [28]. Any force controller may be used; proportional control was used in this reference. The goal of this strategy is to damp out the oscillations caused by the transition. While this may be successful for soft environments, stiff environments have oscillations with small amplitudes and high frequencies. This makes damping difficult for three reasons. First, changes in position of the environmental surface may be smaller than the resolution the manipulator’s position measurement devices. In this case, no velocity will be sensed. Second, for fast oscillations the calculated velocity signal will lag its ideal value, and the damping force may cause instability by being applied out of phase with the true velocity of the surface [56]. Third, flexion in the links due to impact can slightly change the arm structure, thereby making the kinematics and velocity signal computation erroneous. These problems are compounded by the fact that a stiff environment which causes them will also cause a larger impact force and need stabilizing compensation all the more. Thus, this scheme may fail when most needed.

Passive Compliance and Damping Another method for absorbing the shock of impact is to use passive compliance, either on the end effector or in the environment. Some researchers have proposed the use of soft force sensors or compliant ‘skin’ covering for the force sensor [48, 72, 3, 1]. These methods appear to provide stable impact in two ways. First, the material used naturally provides passive damping that helps absorb some of the energy of impact, without the resolution or time lag problems of active damping. Second, the compliance of the material effectively lessens the stiffness of the system composed of the material and the environment. Following from the argument of the previous paragraph, this lessening of the stiffness helps active damping work. Because the end effector remains in contact with the environment over a larger range of displacement for the same experienced force, the displacement will not be below the resolution of the arm’s position (and therefore velocity) measurement devices. Also, the frequency of oscillation will be less, reducing the phase lag of the computed velocity which is needed for active damping.

But there are problems with passive compliance. First, it may not be modified without physical replacement of the material. Second, it limits the effective stiffness of the manipulator during position control. Third, it eliminates precise knowledge of the position of the environment. And fourth, it limits the forces that may be applied — beyond a certain range of operation the compliant material is not linear and is prone to physical failure.

Integral Explicit Force Control Integral force control acts as a low pass filter [76, 64, 55]. Thus, for impact transients, the high frequency components are filtered effectively. For impacts with low energy or with an inelastic environment this may be sufficient [76]. Otherwise, bouncing and possible instability will occur [65]. This is because of the nonlinear loss of contact with the surface and subsequent integrator wind-up which cause severe hopping on the surface.

Impedance Control and Proportional Explicit Force Control It has been analytically and experimentally demonstrated that impedance control against a stiff environment is equivalent to proportional gain explicit force control with feedforward [56, 61, 65]. While both schemes have been tried by many researchers, the gains in these implementations are typically not tuned for the best impact response [28, 75, 25, 21, 3]. For explicit force controllers the gain is tuned for optimal command following once contact has been established. Equivalently, the mass ratio of impedance control is chosen to obtain the desired inertia for free space motion or force exertion, but not impact. The result is an oscillatory system in which bouncing occurs after impact. This is consistent with simulation and experimental results [15, 3, 65]. A solution to this problem is to use a different proportional gain for the impact phase. To understand the proper choice for the gain values it is necessary to analyze both explicit force control and impedance control schemes with a proper system model. This model is reviewed next, followed by a review of the force control strategies.

4 Arm / Sensor / Environment Model

The physical system employed in the study of the robot impacting the environment is depicted in Figure 8. Note that the arm and the environment are part of the same kinematic chain, since they are both attached to mechanical ground. For many of our experimental tests, environment is a cardboard box with an aluminum plate resting on it. The measured stiffness of this environment is $\sim 10^4$ N/m. The box is resting on a table that is considerably more stiff than the box, and is therefore considered ground for these tests. The force sensor is mounted on link six of the manipulator, the CMU DD Arm II. It has a measured stiffness of 5×10^6 N/m. Attached to the force sensor is a steel probe with a brass weight on its end. The brass weight serves as an end effector substitute and provides a flat stiff surface for applying forces on the environment. Previous analysis has indicated that a fourth order model of this arm / sensor / environment is necessary and sufficient for force control. This section presents a review of the development of this model. Full details may be found in [14, 56, 59].

The dynamics of an n DOF, serial link manipulator are described by a set of nonlinear, coupled differential equations [17]. Included in this description are Coriolis and centripetal forces as well as viscous damping and gravitational loading. However, these elements of the description may not always be significant. For instance, Coriolis and centripetal forces are not present when the manipulator is statically exerting force on the environment; viscous damping is not present in direct drive motors; and gravitational loading is not present for a space based robot. Further, active compensation can remove the torques due to these physical effects. For instance, calculation of the inverse dynamics of the arm removes the effects of gravity loading and Coriolis and centripetal forces [29], and negative damping gains can remove the effects of viscous friction [11].

Therefore, for the purpose of this discussion the most important component of the dynamic description of the manipulator is its inertia. All other nonlinear components of the description will be ignored, and assumed to be insignificant or compensated for. Further, by the appropriate transformation the description of the manipulator dynamics may be represented in Cartesian

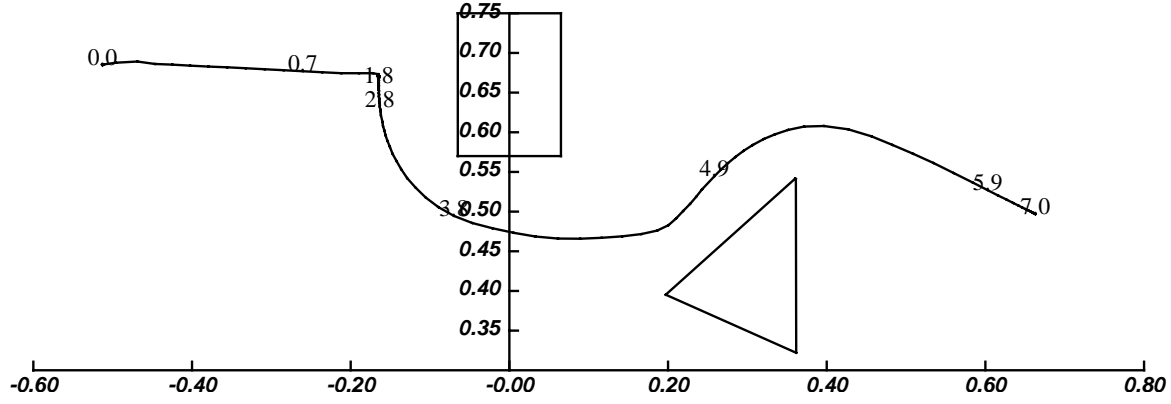


Figure 7: Experimental data showing a path taken to successfully navigate between two objects. For the rectangle, the potential parameters are: $\alpha = 6$, $A = 0.1$. For the triangle, the potential parameters are: $\alpha = 3$, $A = 0.01$. The numbers along the paths indicate the time in seconds.

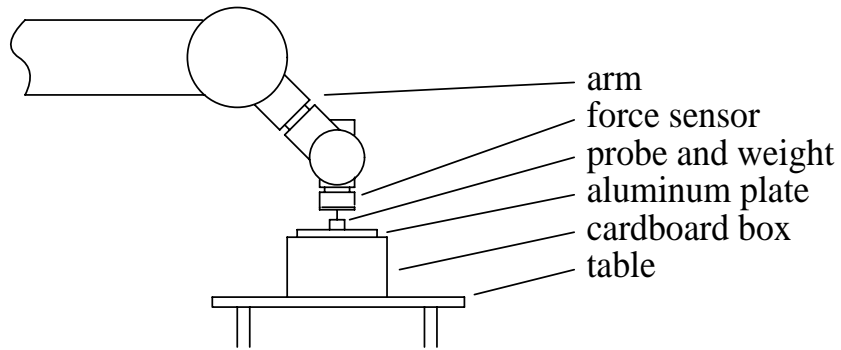


Figure 8: Experimental setup for force control experiments.

space [28]. If the task frame in Cartesian space is aligned with the principle axes of the inertia tensor, the dynamic description becomes fully decoupled. In this case, a single degree of freedom may be considered independently.

The most basic one DOF model of a manipulator is a second order model that has a single mass, damping, and stiffness for the manipulator. The mass is configuration dependent and represents the effective manipulator inertia in that degree of freedom. The damping, if it exists, is a combination of the projection of the viscous joint damping into Cartesian space, and the active damping which may be performed directly in Cartesian space. The stiffness is due to the combination of mechanical and actively applied stiffnesses. The mechanical stiffness can come from either the links or the actuators. For now, we will ignore the link stiffness and consider the links to be pure transmitters of force. Actuator stiffness typically comes from gearing which is non-backdriveable. Many manipulators are backdriveable and do not exhibit mechanical stiffness. The CMU DD Arm II has no joint friction or gearing and, therefore, damping and stiffness will only be present if provided actively [29].

Having introduced a model of the manipulator, it is necessary to discuss an environmental model. Some researchers have made no assumptions about the structure of the environment, and have assumed instead that interaction with it will produce measurable forces [49, 22, 20, 36, 26, 18]. Other researchers, usually those working with a compliant system or sensor, have modelled the environment as a mechanical ground [72, 51]. Still others, have recognized that the environment has some compliance, and therefore have modelled it as a simple stiffness [69, 45, 46, 70, 28, 12, 75, 13, 24, 32, 23]. Finally, some researchers have modelled the environment as a complete second order system with components of mass and damping, as well as stiffness [26, 14, 15, 76]. This last form of the environmental model recognizes that the environment has oscillatory modes of its own, but simplifies the overall analysis by only considering the first mode. Thus, the second order model is more restrictive than just a general environment that exerts measurable force on the arm. But the specific representation of the model's dynamic components will permit a better understanding of the interaction between the arm and the environment.

Between the arm and the environment exists the force sensor. While a very stiff force sensor may not always exhibit its dynamics, under certain circumstances they may become important. The use of a stiff robot position controller, contact with a stiff environment, or impact with the environment may excite the sensor dynamics. Therefore, it is sometimes necessary to include the sensor in our model. A second order model of the sensor dynamics can be added to the above models of the arm and environment by placing a spring and damper between the masses of these two second order systems.

Finally, it is necessary to return to the subject of link stiffness and higher order arm dynamics. It has been recognized by some researchers that the arm has higher order dynamics that may need to be modelled [48, 14, 15, 76]. This is particularly true if the environment and sensor are stiff. Inclusion of a second order approximation for the link stiffness makes the composite arm model fourth order, and the arm / sensor / environment model sixth order. However, if the link and sensor dynamic characteristics are similar, they may be lumped together. For instance, a typical force sensor is composed of strain gauges mounted on aluminum. If such a sensor is mounted on an aluminum robot arm, there is no clear distinction between the end of the last link and the beginning of the sensor. Modelling just the first mode of vibration of this entire assembly requires only a second order model for both the arm links and the force sensor. (This concatenation of the stiffness and damping of both components reduces the total stiffness and damping of the link-sensor by a geometric proportionality factor that depends on the arm and sensor designs [56].) Thus the entire model for the arm-actuator / arm-linkage-and-sensor / environment system can be reduced from sixth to fourth order. This model is shown in Figure 9. The transfer function of this system is:

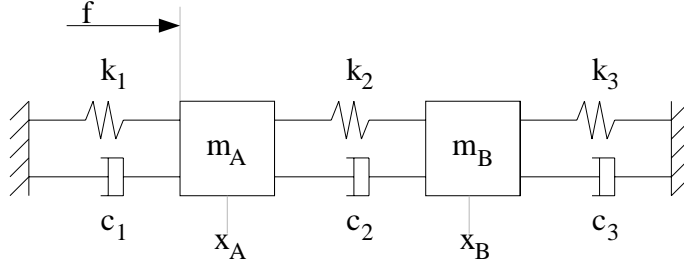


Figure 9: General fourth order model of the arm, sensor, and environment system.

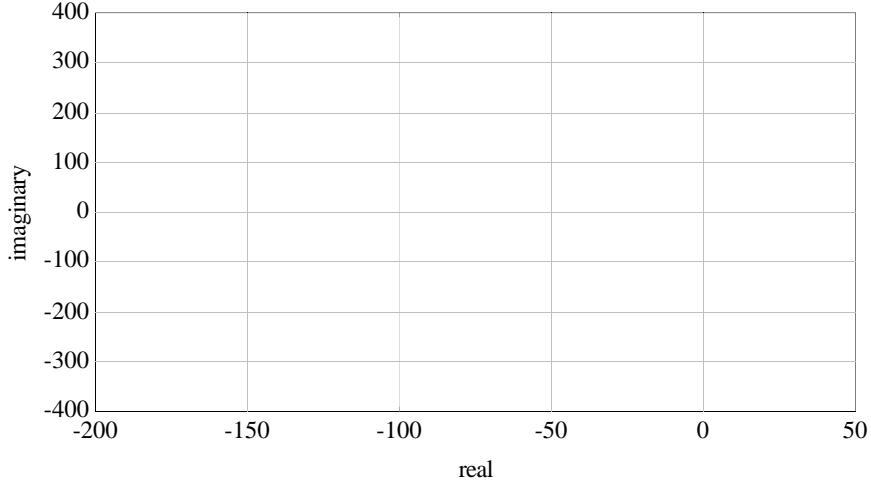


Figure 10: The pole and zero locations for the fourth order model, using the experimentally extracted parameters. Not shown is the leftmost pole which is at -28000 on the real axis.

$$\frac{F_m}{F} = \frac{(m_B s^2 + c_3 s + k_3)k_2}{(m_B s^2 + (c_2 + c_3)s + (k_2 + k_3))(m_A s^2 + c_1 s + k_1) + (m_B s^2 + c_3 s + k_3)(c_2 s + k_2)} \quad (5)$$

where $F_m = k_2(X_B - X_A)$ is the measure force; x_A is the measured position of the arm; x_B is the position of the environment; and m , k , and c are the mass, stiffness, and damping parameters of the fourth order model. This is similar to the model presented in [15].

We have experimentally extracted parameter values for the components of this model for the described experimental configuration. Theoretical and experimental details can be found in [59]. The pole/zero locations indicated by the extracted values differ greatly from those assumed by other researchers [14, 15]. Figure 10 shows all but the leftmost pole, which is at -28000 on the real axis. The complex pole pair is due mainly to the environment. The real pole pair (the real pole shown plus the other not shown) is due mainly to the sensor dynamics. These pole pairs will be called the environment and sensor poles, respectively. It can be seen that the sensor poles are fairly far removed from the environmental ones, and are located farther into the left half plane. The leftmost sensor pole will be ignored since it is negative real, and far removed from the others.

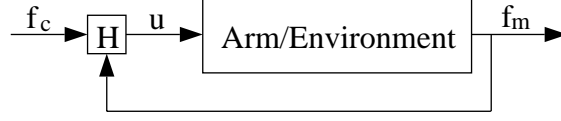


Figure 11: Explicit force control block diagram.

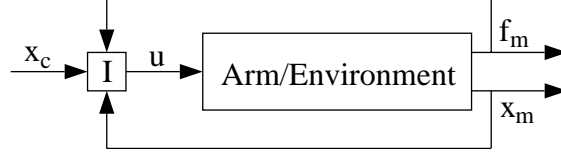


Figure 12: Impedance control block diagram.

5 Explicit Force Control and Impedance Control

The system modelled in the previous section is the plant of the controller used for environmental interaction. Two main conceptual choices have emerged for the choice of this controller structure: *explicit force control* and *impedance control*. It has been shown both theoretically and experimentally that second order impedance control against a stiff environment is essentially equivalent to proportional gain explicit force control with feedforward [56, 61]. The argument supporting this conclusion will only be reviewed here.

First, it is necessary to present the block diagrams of the explicit force and impedance controllers, as in Figures 11 and 12. Next, it is important to recognize that the linear impedance controller may be separated into a position component and a force component, as in Figure 13. Further, Figure 14 shows that because there is no external reference force signal, the force loop may be considered an internal explicit force controller. The type of this internal explicit force controller can be extracted by looking at the impedance control law [20, 56]:

$$\tau = J^T \Lambda M^{-1} [C(\dot{x}_c - \dot{x}_m) + K(x_c - x_m) - f_m] - J^T \Lambda \dot{J} \dot{\theta} + h + g + J^T f_m \quad (6)$$

where Λ is the manipulator inertia matrix in Cartesian space; M , C , and K are the second order impedance matrices; h is the vector of Coriolis and centripetal forces; g is the gravitation force vector; J is the manipulator Jacobian; f , τ , x , and \dot{x} are vectors of force, torque, position, and velocity; and subscripts c and m indicate commanded and measured quantities. The terms that compensate for velocity dependent forces and gravity can be considered feedforward terms, and ignored for the remainder of this discussion. What is left is an equation for torque of the form:

$$\tau = J^T [H'(f_c - f_m) + f_m - K_v \dot{x}_m] \quad (7)$$

$$f_c = K(x_c - x_m) + C \dot{x}_c \quad (8)$$

$$H' = \Lambda M^{-1} \quad (9)$$

$$K_v = H' C \quad (10)$$

The active damping provided by K_v may be added to the passive damping in the plant (c_1 in Equation (5)) and removed from further consideration in the control equations.

Thus, the internal explicit force controller in impedance control can be represented by the block diagram in Figure 15, where G is the plant given by Equation (5). In this figure, the positive

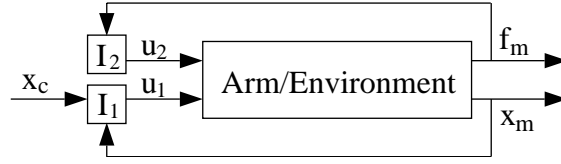


Figure 13: Impedance control block diagram with the controller divided into its position part, I_1 , and its force part, I_2 .

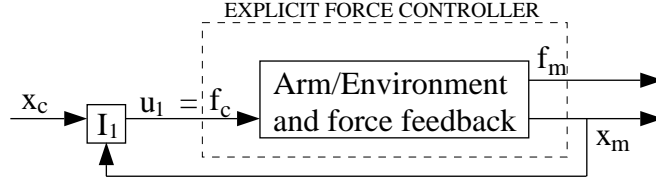


Figure 14: Impedance control block diagram redrawn to show the inner explicit force controller.

feedback loop acts as a reaction force compensation. If the sensor dynamics are ignored, the physical reaction force loop may be directly extracted from the plant [56]. As seen in Figure 16, this creates a new plant, G' , and a negative feedback loop of the physical reaction force. Further, this figure shows an equivalent expression of the proportional gain as $H' = H + 1$. The transfer function for this controller is:

$$\frac{F_m}{F_c} = \frac{H'G'}{1 + H'G'} \quad (11)$$

$$= \frac{(H + 1)G'}{1 + (H + 1)G'} \quad (12)$$

It can be seen directly that an equivalent block diagram of this system may be constructed as in Figure 17. This is a proportional gain explicit force controller with feedforward force and serves as the inner force loop in the impedance controller.

Therefore, the impedance controller has the form of a PD position controller surrounding a proportional force controller. But when in contact with a stiff environment, the position of the environment can be set as the origin ($x_m = 0$). Also, the commanded velocity is usually zero ($\dot{x}_c = 0$). Thus, the external position loop of the impedance controller provides a command force that is simply: $f_c = Kx_c$. The external position loop becomes, in effect, functionless. We conclude that *for the case of a stiff environment, impedance control is equivalent to proportional gain explicit force control with feedforward*. Experimentation has validated this conclusion [61].

It is interesting to look at what this equivalence implies for gain value selection. (In this discussion only the one dimensional or diagonal matrix case will be considered.) First, the stability of the impedance controller is guaranteed for $H' \geq 0$. This is equivalent to the condition $\Lambda M^{-1} \geq 0$. Assuming a constant manipulator inertia Λ , gain H' varies as the inverse of the target impedance mass, M . Zero gain means infinite mass, and large gain means small mass. For the proportional force controller, stable gain values are $H \geq -1$ since $H = H' - 1 = \Lambda M^{-1} - 1$. *Negative proportional force control gains down to minus one are stable*. Further, it will be seen in the next section, that they are desirable for impact control.

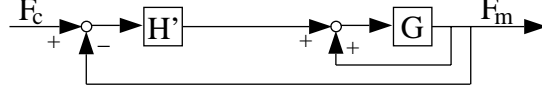


Figure 15: Block diagram of a force-based explicit force controller with proportional gain and positive feedback for reaction force compensation.

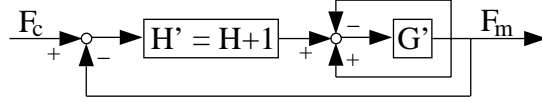


Figure 16: Block diagram of a force-based explicit force controller with proportional gain and extra feedback for reaction force compensation. The plant G has been expanded into its components, and the sensor dynamics have been ignored.

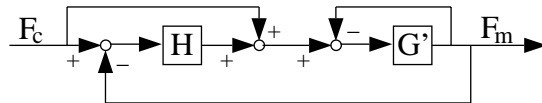


Figure 17: Block diagram of a force-based explicit force controller with proportional gain and unity feedforward. The plant G has been expanded into its components, and the sensor dynamics have been ignored.

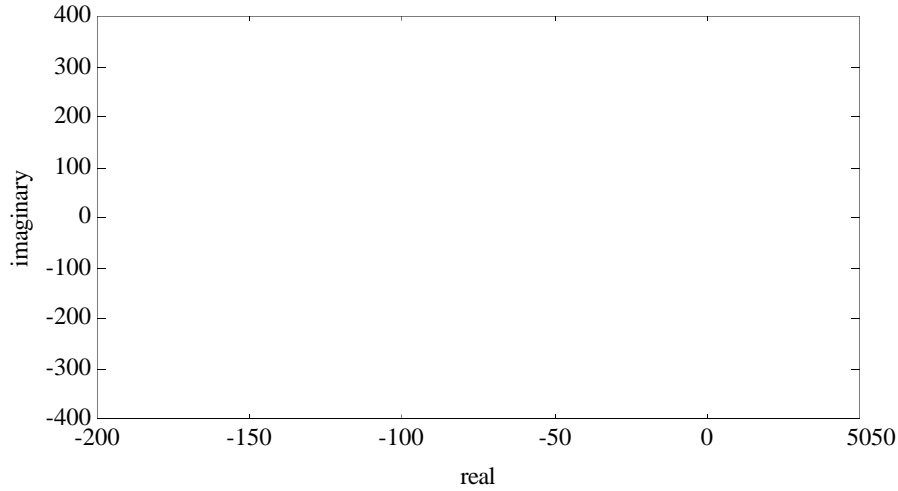


Figure 18: Root locus for the second order model for $H' = H + 1$. The double root occurs for $H' \approx 1.5 \times 10^{-3}$. The poles shown in the middle of the locus are for $H = 0$ and correspond to the environmental pole locations in Figure 10.

5.1 Impact Control Without Sensor Dynamics

The model of the arm / environment plant that neglects sensor dynamics results in pole and zero locations similar to those shown in Figure 10, except the sensor poles are not present. For a proportional gain explicit force controller with this plant, the root locus is shown in Figure 18 ($H \geq -1$). (The poles shown in the middle of the locus are for $H = 0$ and correspond to the environmental pole locations in Figure 10.) Note that one pole will go into the right half plane for $H < -1$ as predicted. Observing this root locus it is immediately apparent that the most stable gain is the one that places the two poles at the point where the roots leave the real axis. Ignoring the sensor dynamics, an approximate value of this gain may easily be determined [56]. The double root of the characteristic equation occurs for a value of the proportional gain close to negative one. There are three equivalent ways to view or interpret this result:

Proportional force control with reaction force compensation. This is the controller in Figure 16. In this case, the controller does not utilize the force error signal since $H' \approx 0$. However, the reaction force of the impact is directly negated by a feedback signal. Viewed this way, the impact controller does not bounce because the oscillations in the commanded force and those in the experienced force are equal and opposite. Thus the surface is at a node of two interfering pressure waves. No net force means no net acceleration. Any initial oscillation is damped out by natural and active damping.

Proportional force control with negative gain and a feedforward signal. This is the controller in Figure 17. While this controller looks different than above, it has been shown previously that it is equivalent. In this case the controller multiplies the force error by $H = H' - 1 \approx -1$. There is also a feedforward signal of the commanded force.

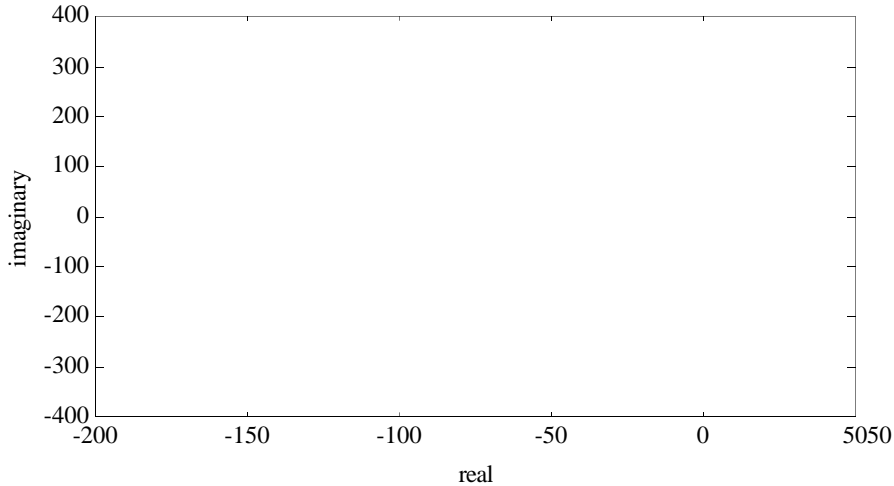


Figure 19: Root locus for the fourth order model for $-1 \leq H < \infty$ or $0 \leq H' < \infty$. The poles shown in the middle of the locus are for $H = 0$ and correspond to the environmental pole locations in Figure 10.

Impedance controller with a large target mass. As discussed previously, an impedance controller is equivalent to an explicit force controller when in contact with a stiff environment. Impedance controllers employ a proportional gain, ΛM^{-1} , where Λ is the arm inertia, and M is the desired inertia. Viewed in this way, the impact controller matches the apparent mass of the arm to the stiffness and damping of the environment such that the resultant system is critically damped. More imprecisely, it can be said that the arm is made to appear so massive that it can't bounce.

5.2 Impact Control With Sensor Dynamics

Including the sensor dynamics changes the above analysis somewhat by introducing an additional set of poles. Obviously, if the sensor poles are far from the environmental poles they will have little effect, and the above results will remain the same. However, the fourth order model that was previously developed has one pole relatively close to the environmental poles and zeros. Figure 19 shows the root locus for this system for proportional gain values of $-1 \leq H < \infty$ or $0 \leq H' < \infty$. The points of closest approach of the locus to the real axis correspond to gain values of $H \approx -0.8$ or $H' \approx 0.2$. These are the best values for impact control.

It is important to point out that this locus also indicates that positive gain proportional force control, as well as impedance control, will become unstable for this system. The instability of these schemes has been confirmed experimentally [56, 61, 63, 64]. The points on the locus in the right half plane correspond to values of $H > 1$ or $H' > 2$, and have been shown to be unstable. While the root locus suggests that very high gains would again be stable, experimentation has indicated that the system model breaks down for these large parameter values.

5.3 Impact Experimental Results

Figures 20 show the results of impact tests on the modelled environment using impact control. The solid line is the measured force; the dashed line is the reference force, which is non-zero only after impact; and the dotted line is the scaled measured velocity. As can be seen, the impacts occur at a

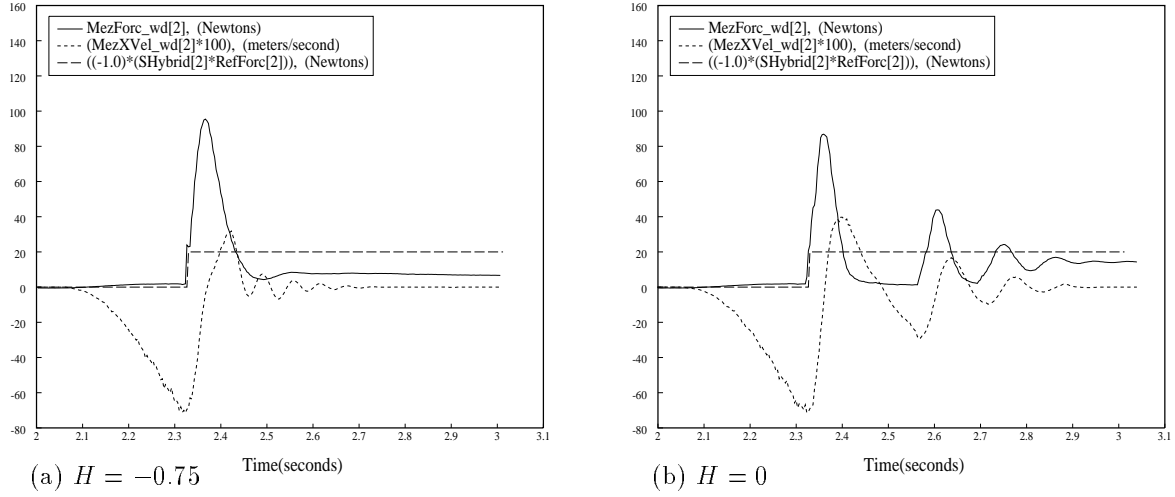


Figure 20: Experimental data comparing the response of impacts on the test environment using proportional force control with feedforward force. Figure (a) shows the result of using a negative proportional gain of -0.75. Figure (b) shows the open loop response to the impact.

speed of 0.75 m/s. Figure 20(a) shows the stability provided by negative proportional gains. This is in contrast to the open-loop response shown in Figure 20(b). For positive gains the bouncing becomes more severe, until instability results [65].

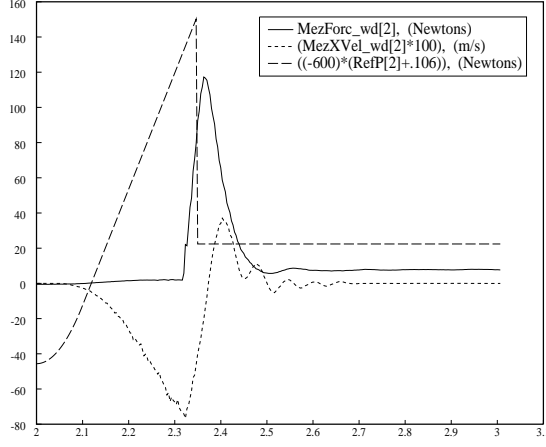
Figures 21 show the equivalent response of impedance control to the impact transient, as predicted. In these figures, the solid line is measured force; the dashed line before the impact is commanded position; the dashed line after impact is the commanded position minus the environment position, multiplied by the impedance stiffness parameter; and the dotted line is the measured velocity. As predicted, Figures 20 and 21 show impedance control provides response equivalent to proportional gain control with feedforward force.

To further test the impact controller, collisions with a very stiff steel environment were performed [65]. The results of these tests are shown in Figure 22. Again, the impact controller eliminates hopping and provides stability.

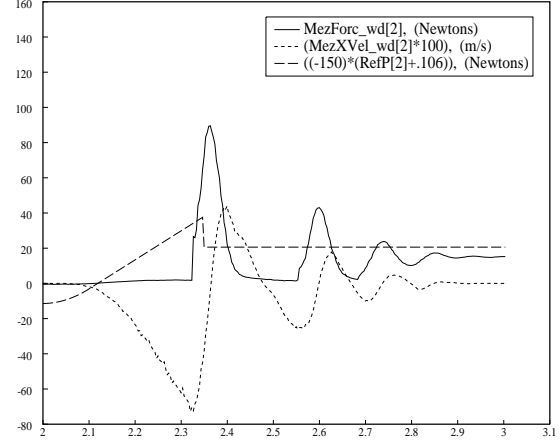
6 Explicit Force Control

After impact control has successfully provided stable contact transition for the manipulator, it is desirable to provide accurate force control for the system. This section will review commonly proposed techniques and experimental data evaluating them. Also, since its equivalence to proportional gain force control has already been shown, impedance control will not be reviewed further here.

Explicit force control strategies utilize direct evaluation of desired and measured forces to determine the control signal. Many basic forms of this type of controller have been proposed, and this section will review several of them [64]. Two types of output are possible from force controllers: forces or positions. In the former, forces are commanded directly, and then translated into manipulator joint torques. In the latter, position setpoints are given to an inner-loop position controller.

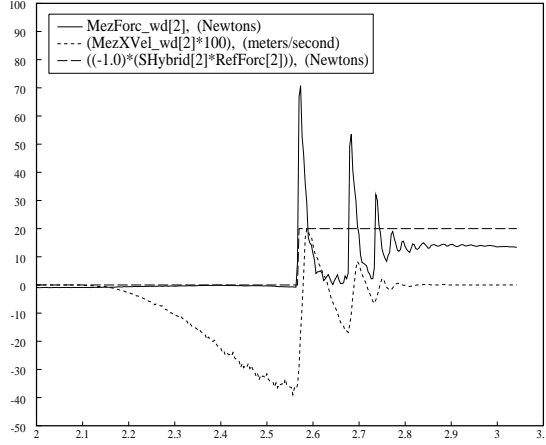


(a) $H = -0.75$ Time(seconds)

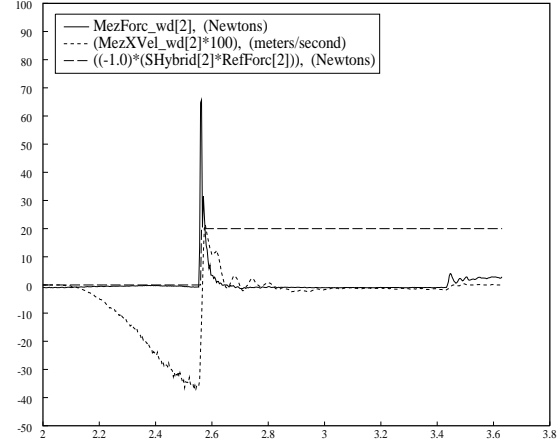


(b) $H = 0$ Time(seconds)

Figure 21: Experimental data comparing the response of impacts on the test environment using impedance control, with mass ratios of 0.25 and 1.0. The response is the same as that shown in Figure 20, indicating the equivalence of impedance control with proportional gain force control.



(a) $H = 0$ Time(seconds)



(b) $H = -0.85$ Time(seconds)

Figure 22: Impact control on the steel pedestal in the z direction: comparison of $H = 0$ (open loop) and $H = -0.85$.

We have previously shown that position-based methods may be recast as force-based methods [62]. Therefore, only the force-based methods are reviewed here.

Force-based explicit force control describes a method that compares the reference and measured force signals, processes them, and provides an actuation signal directly to the plant. The reference force may also be fedforward and added to the signal going to the plant, described in Section 4. To control this plant some subset of PID control (i.e. P, I, PD, etc.) is usually chosen. From a computational perspective, all are approximately equal in complexity. The strategies presented here are generalizations of schemes previously proposed, as indicated below. In all cases, the joint torques commanded by these schemes are obtained through the transpose of the Jacobian, and gravity compensation is employed for terrestrial operation. The parameters f and \dot{x} are Cartesian force and velocity. K is a gain for either proportional (subscript fp), integral (fi), or derivative (fd) force control. K_v is the velocity gain and provides active damping that is incorporated directly into the plant [62]. The subscripts c and m denote commanded and measured quantities. The variables s and a are Laplace domain complex numbers.

Proportional Control [45, 28, 3, 14, 15, 75, 68]

$$f = f_c + K_{fp}(f_c - f_m) - K_v \dot{x}_m \quad (13)$$

Integral Control [52, 76, 10, 55]

$$f = K_{fi} \int (f_c - f_m) dt - K_v \dot{x}_m \quad (14)$$

Proportional–Integral Control [46, 37, 15]

$$f = K_{fp}(f_c - f_m) + K_{fi} \int (f_c - f_m) dt - K_v \dot{x}_m \quad (15)$$

Proportional–Derivative Control [15, 72, 55]

$$\text{Unfiltered : } f = f_c + K_{fp}(f_c - f_m) + K_{fd} \frac{d}{dt}(f_c - f_m) - K_v \dot{x}_m \quad (16)$$

$$\text{Filtered : } F(s) = F_c(s) + [K_{fp} + K_{fd}s] \left[F_c(s) - \left(\frac{a}{s+a} \right) F_m(s) \right] - K_v s X_m(s) \quad (17)$$

Second Order Low Pass Filter Control [56, 62]

$$F(s) = \frac{K_{fp}}{s(s+a)} (F_c(s) - F_m(s)) + K_v s X_m(s) \quad (18)$$

6.1 Experimental Evaluation of Force Control

Of the proposed force control strategies, integral force control has proven superior to the others [64]. This is mainly due to the noisy nature of force signals, which the low-pass integral controller effectively filters. The proportional controller is not strongly effected by the noise, but the PD controller is driven unstable by it. (As will be seen later, even filtering could not improve the response of the PD controller.) Also, the best response of the second-order low pass filter is when one pole dominates and the behavior is like the integral controller.

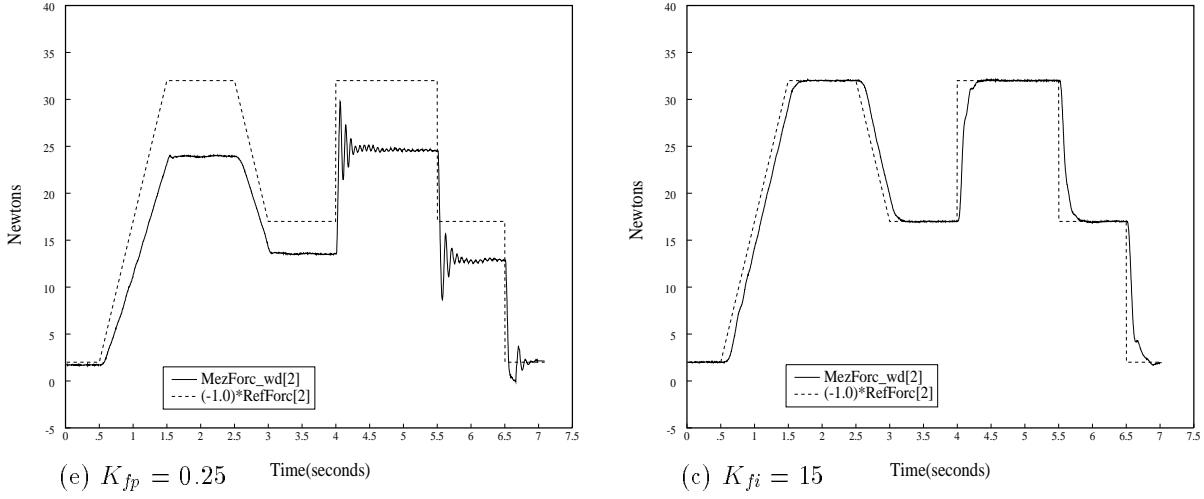


Figure 23: Experimental data of explicit force control schemes. These are the best results obtained for (a) proportional control with force feedforward and (b) integral gain control.

Therefore, the two main force controllers of interest have proven to be the proportional and integral gain controllers. Figures 23 show the best responses obtained with each controller in contact with previously discussed environment. The solid and dashed lines are the measured and reference forces, respectively. Note that the proportional controller has a steady-state error, which can not be decreased without increasing overshoot and instability [64]. Alternatively, the integral controller has lag, but follows the form of the reference and has no steady state error. Finally, the response of the PI controller is intermediate to these two controllers and remains inferior to the integral controller.

6.2 Blending Impact and Force Control

The transition from position control to impact control is abrupt, and triggered by the impact force spike. But the transition method from impact control to integral force control is less obvious. One method of doing this is to have a transition period in which the proportional gain and force feedforward of the impact controller are brought to zero, while the integral gain is increased to its best value. We propose a linear switch of all gain values after the impact force pulse diminishes. Determination of the times to begin and end the transition is currently obtained empirically. The beginning and ending gain values are determined from tests of the impact and integral gain force controllers interacting with the same environment [63].

Figure 24 shows the results of this strategy. The impact control phase lasts for 0.15s (about the width of the impact spike) after the beginning of the impact. This is followed by a period of transition from impact control to integral gain force control which lasts 0.15s. During the transition phase H is varied linearly from -0.75 to 0; K_{ff} , the feedforward gain, is varied linearly from 1 to 0; and K_{fi} is varied linearly from 0 to 15. After, this transition period, integral force control is continued with the gain at fifteen. As can be seen, this simple strategy provides stability through the impact period and excellent position and force control before and after the impact.

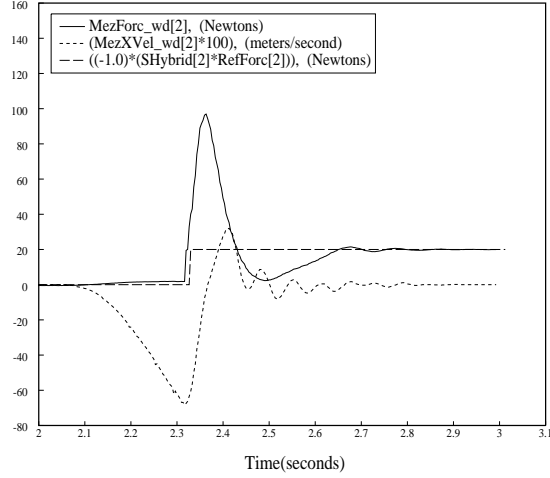


Figure 24: Experimental data of impact control with transition to integral gain force control. The impact control phase lasts for 0.15s after the beginning of the impact. This is followed by a period of transition from impact control to integral gain force control which lasts 0.15s. Beyond 0.3s after impact, integral gain force control is used.

7 Algorithm Implementation Considerations

In addition to the above control issues, there are many problems and issues associated with the implementation of the previously reviewed algorithms. Some are only minor annoyances, while others can effect the stability or range of operation of a particular controller. Among the issues discussed in this section are position measurement resolution, velocity signal calculation, force signal noise and filtering, the force signal derivative, hybrid control switching, and impact transient handling.

Much of the following discussion utilizes graphed data to illustrate and validate the topic. In the legends of the graphs, the reference value of the applied force is called ‘RefForc’; the measured value of the experienced force is called ‘MezForc_wd’; the filtered value of this measured force is called ‘Filter_Forc’; the derivative of this filtered force is called ‘Fdot’; the measured value of the Cartesian velocity is called ‘MezXVel_wd’; the measured value of the end effector position is call ‘MezP’; and the hybrid control selection parameter is called ‘SHybrid’ [46]. All of these variables are vectors and the indices follow the conventions of the C computer language.

7.1 The Position Signal

The position signal is a 16 bit absolute position value obtained from pancake resolvers located at each joint. Therefore the resolution of each joint is $\approx 10^{-4}$ radians. Since the CMU DD Arm II has a reach of approximately one meter, the position error incurred from all six joints is about one millimeter. Therefore, the workspace of the manipulator can be thought of as quantized into small volumes, a millimeter in diameter. Obviously, the mapping from joint space error to task space error does not generate consistently shaped volumes, but the more important conceptual fact of quantized position measurements remains valid.

This intrinsic position error is important for both obstacle avoidance and force control, for the same reason. Both the physical world, and the computed values of the potential functions,

are continuous at the millimeter scale. But the measured positions obtained from the arm are effectively discontinuous at this scale. Therefore, if forces from the potential function or physical environment change drastically from one millimeter volume to the next, then instability is likely. Typically this will appear as chattering at the boundary between zero force and a large force. This phenomenon can occur for controllers in contact with stiff environments, or steep potential functions. In both cases, this extra constraint is essentially a step function to the arm controller. If the environment is made less stiff, or the potential function less steep, then the spatial quantization effects will be reduced, and the chattering will disappear. When implementing the superquadric potential functions, this criterion places a qualitative upper bound on the value of α , which effects the potential function steepness.

7.2 The Velocity Signal

The angular velocity signal for the joints of the CMU DD Arm II is obtained by differencing and averaging the angular position signal. Every control cycle, the position is obtained and placed in a stack. A velocity signal averaged over the past n control cycles can be obtained by simple differencing of the current position with the one n cycles before it:

$$\begin{aligned} v_{avg} &= \frac{1}{n} [v(t) + v(t - T) + \dots + v(t - nT)] \\ &= \frac{1}{n} \left[\frac{p(t) - p(t - nT)}{T} + \frac{p(t - T) - p(t - 2T)}{T} + \dots + \frac{p(t - (n - 1)T) - p(t - nT)}{T} \right] \\ &= \frac{1}{nT} [p(t) - p(t - nT)] \end{aligned} \tag{19}$$

Good results are obtained for the CMU DD Arm II with $3 \leq n \leq 10$. The lower number provides a velocity signal with less lag and more noise, and the higher number just the opposite. For free space motion with the CMU DD Arm II, the natural frequency of the system is determined by the stiffness provided by the position gain. This frequency is usually low enough that the velocity signal lag is not significant. However, when the arm is in contact with the environment, the natural frequency of the system is largely determined by the environmental stiffness. This frequency is much higher than in the free space motion case. Therefore, the velocity signal lag can become a major portion of the oscillation cycle. As the delay approaches 90° the velocity signal will be in phase with the position signal. In this case, the velocity gain will not damp, but rather add to the already large stiffness of the system, driving it toward instability.

For the tests conducted in contact with the environment, a velocity averaging factor of $n = 3$ was used for joints 4, 5, and 6. A factor of $n = 5$ was used for joints 1, 2, and 3. For tests involving free space motion, the natural frequency is smaller and a factor of $n = 10$ for joints 1, 2, and 3 is usually used. The value $n = 3$ for the last three joints tends to be sufficient at all times.

Since the Cartesian velocity signal contains components from all of the joint signals, the delay will be between three and five cycles. For the control rate of 300 Hz, the delay is between 0.01 and 0.016 seconds. Figure 25 shows the velocity and position signals during proportional gain explicit force control ($K_{fp} = 0.75$), after a step input. The delay of the velocity signal is about 0.01 seconds, or a 45° phase lag for the 12 Hz oscillation. This also explains why an averaging factor of $n = 10$ is unacceptable. This delay would put the velocity signal in phase with the position signal.

Note that active damping when the manipulator is in contact with the environment must be used with caution. The time delay from the velocity calculation/filtering is always present. If this delay is a significant part of the natural frequency of the system, then the velocity signal will act as a position signal and add to instability. Further, stiffer environments have higher oscillation

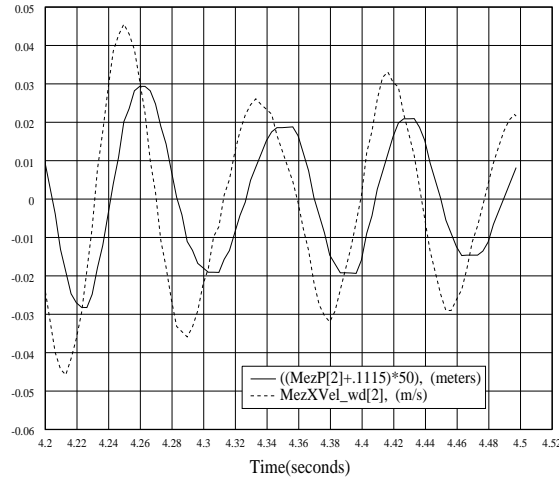


Figure 25: Velocity phase lag due to averaging. It can be seen that the velocity signal lags its ideal value by about 0.16s, or 45 degrees.

frequencies, making the velocity signal least reliable when it would be most useful. Therefore, the damping intrinsic to impedance control, and sometimes used in explicit force control, is always suspect.

Finally, it is worth mentioning that these problems with delay only apply to active damping. Passive damping, as supplied by some soft sensors or end effector covers, will provide damping without time delay [72, 3]. These devices also lower the natural frequency of the system, making active damping possible.

7.3 The Force Signal

A Lord 15-50 force sensor was used in all of the experiments. In its factory configuration, it supplies eight strain gauge values at 416 Hz. However, the controllers used often ran at only 300 Hz, as previously discussed. Since the Lord sensor controller has its own internal clock, there is no way to easily change the update rate. Alternatively, individual request for data can be made, but only with a maximum rate of 250 Hz, due to clock skew. Therefore, we chose to receive the data at the faster rate of 416 Hz and ignore one of every four sets. This has the added effect that each data set could be as old as one 416 Hz cycle, or 2.4 ms. This asynchronous sampling has no appreciable effect on the stability of the controllers since force oscillations were an order of magnitude slower than the control rates and sampling time.

What does drastically effect control stability is the noise in the resultant force signal. As shown in Figure 26, this can be substantial. Filtering of the force signal is partially effective, but introduces lag as can be seen in Figure 27, where the measured force signal is a solid line and the filtered force signal is a short dash line. This effect is very detrimental for PD force control [63], as will be reviewed in the Section 7.4.

The force signal noise has several contributors, discussed below: intrinsic noise, kinematic fluctuations, kinematic inaccuracies, and inertial effects. All of these factors contribute to a noise amplitude of ~ 0.1 Newtons. This is an order of magnitude above the sensor resolution.

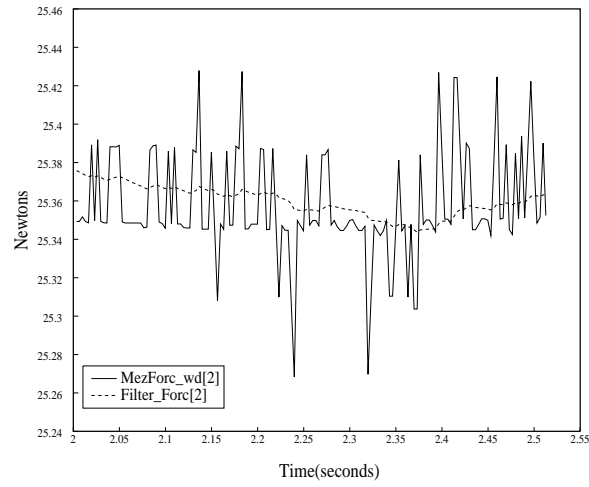


Figure 26: Filtered and unfiltered force signals.

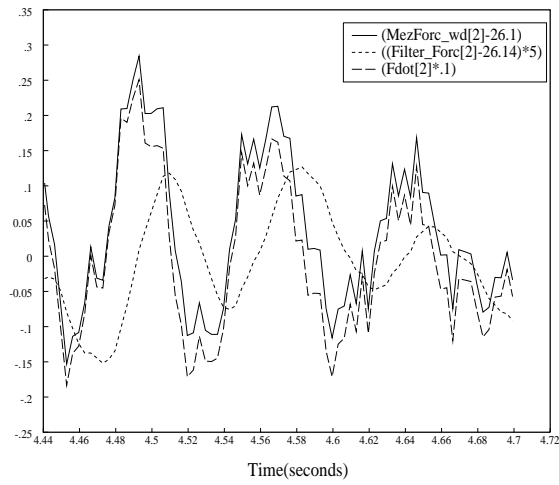


Figure 27: The lag of the filtered force causes the force derivative to be in phase with the measured force.

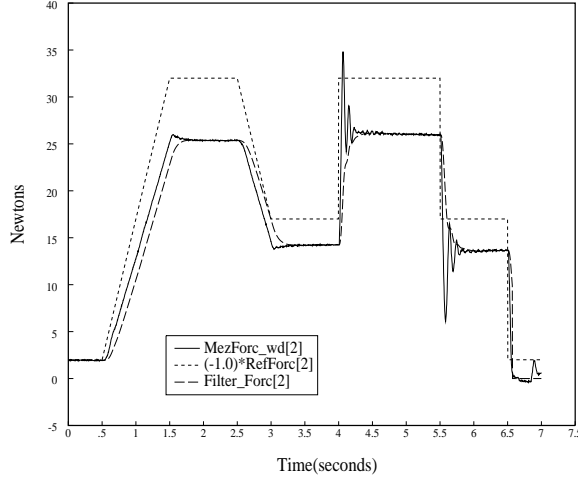


Figure 28: Experimental data from PD control with $K_{fp} = 0.5$ and $K_{fd} = 0.01$.

Intrinsic Noise This is present in the analog and digital electronics of the sensor system as well as the joint position resolvers. The joint position measurement noise contributes because of the need for transformation of the measured force to the control frame. Considering that the CMU DD Arm II has 16-bit absolute positioning resolvers, fluctuation in the last bit typically causes angular errors on the order of $2(2\pi/2^{16}) = 96\mu\text{radians}$. The additive effect of the six joints will then cause a worst case error of 0.57 milliradians in orientation. Multiplying this by the 1.25 meter reach of the arm, gives a position error of 0.71 mm.

Kinematic Fluctuations The kinematics of the arm are based on the assumption that the links are completely rigid. However, bending or oscillations in the arm structure lead to erroneous calculations of the sensor frame position and orientation, and therefore the measured force. This is especially true when there are forces exerted on the arm, such as during impact and force control.

Inertial Effects Inertia causes the measured force not to equal the applied force if the sensor is accelerating. Since most environments are stationary, zero acceleration usually implies that the arm has zero velocity as well.

Many manipulators (including the CMU DD Arm II) are capable of rapid acceleration, and the inertial forces can be considerable. We have observed the inertia of the end effector causing problems with both impedance and explicit force control schemes. For impedance control, the manipulator drifts, since it is attempting to apply an impedance to an external force, when one is not actually present. For explicit force control, the hybrid controller will switch from position to force control in the direction of the inertial force.

7.4 The Derivative of the Measured Force

The previous section described the noise that is present in the force signal. This noise makes it essentially impossible to use PD force control. Even with low-pass filtering, the system was unstable for appreciable derivative gain values. Figure 28 shows the response of the system (solid), as well as the reference force (short dash), and filtered force (long dash), for $K_{fp} = 0.5$, $K_{fd} = 0.01$, $K_v = 10$, and $a = 10$ in Equation (17). The results are not much better than for proportional gain alone [63].

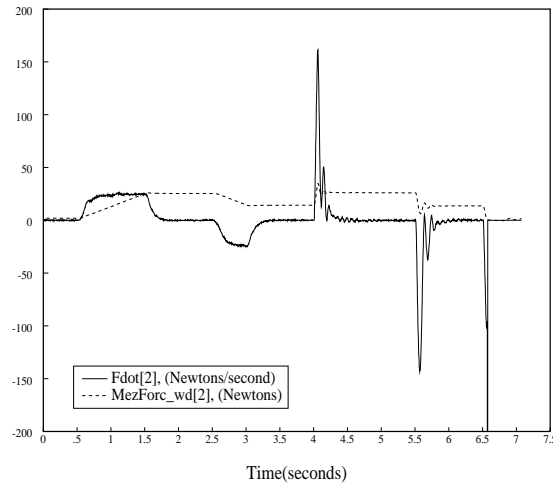


Figure 29: Calculated force derivative and measured force signal used in PD control.

As will be described below, improvements in the performance of this controller can not be made by varying the gains given here.

First, increasing the derivative gain does not improve the response of the system because the amplified low frequency noise can still drive the system unstable. While Figure 28 seems to show a fairly smooth filtered force signal, a close-up view of the same data has already been shown in Figure 26. Much of the noise has been removed, but with a large enough gain the noise will dominate. Making the cut-off frequency of the filter lower ($a < 10$) will eliminate this noise, but it introduces a more serious problem of lag.

Figure 29 shows that the calculated derivative (solid curve) appears accurate. (The dotted curve is the measured force.) However, it is apparent from this figure and Figure 28 that there is lag introduced by the filtering process. This lag becomes extremely important when it is a significant portion of the period of oscillation of the system. Figure 27 shows the original force signal (solid), the filtered force signal (short dash), and the derivative of the filtered signal (long dash). For this oscillation frequency, the filtering process causes the filtered force to lag the measured force by one quarter cycle. This makes the force signal 180° out of phase with the ideal derivative signal. Thus, the proportional gain acts as a destabilizing negative derivative gain. Further, the derivative of the filtered signal leads it by one quarter cycle. Thus, the derivative is in phase with the originally measured force and the derivative gain acts as a proportional gain. Increasing the derivative gain causes greater oscillations exactly when the effective damping is being reduced by the proportional gain. This obviously will cause the system to go unstable.

It can be concluded from this discussion that the filter pole should be significantly larger than the natural frequency of the system. However, it also must be small enough to effectively filter the noise of the force sensor. These two criteria could not be met with our system. To be fair, most systems will never meet this criteria. Force controlled systems are most challenged by stiff environments that have high natural frequencies. It is unlikely that a sensor can be built that has noise only at frequencies much greater than the natural frequencies of these environments.

One solution, however, is to use a soft force sensor or compliant covering on the sensor. The compliance acts as a low-pass filter with no time delay. In this way, the derivative of the force signal may be used under the condition that the time necessary to calculate it is not significant. In this case, without a noisy force signal, simple differencing of the current and most recent force samples

will usually suffice. Thus, all that is required is that the force sampling frequency is not of the same order of magnitude as the natural frequency of the system. Successful PD force control with a soft force sensor has been reported elsewhere [72]. Soft sensors, however, have other drawbacks such as smaller operation range, mechanical fatigue, unactuated degrees of freedom, nonlinearities, etc.

7.5 Hybrid Control Switching

For an explicit force controller it is necessary to switch from position to force control when using a Hybrid Control framework. One way to achieve this is to switch to force control when a measured force threshold is exceeded. To prevent force signal noise from causing the switch, a value of 2 N was used for the threshold value for switching *to* force control. Also, since the noise still exists while in force control mode, the measured force may drop below 2 N inadvertently. Thus, a lower threshold value of 1 N was chosen for switching *from* force control. The switching strategy was implemented in a joystick process running at 150 Hz. Because the switching was done by the joystick controller, the joystick values could be interpreted as commanded velocity (free space motion), or commanded force (constrained motion).

Another aspect of the switching strategy is that it could be made unidirectional — permitting only switching *to* force control. When unidirectional, force control will remain in effect even if the measured force is reduced below the threshold, as when the manipulator leaves the surface. The behavior of the controller for this case of contact loss can be quite interesting and illustrative [60, 63]. Experience showed that some of the controllers tested were sure to become unstable when surface contact was lost. To prevent damage to the system, bidirectional switching was usually used. (If the end effector lost contact with the environment, the controller reverted to position mode, as can be seen at the tail the data in Figure 28.) To prevent the manipulator from losing contact with the environment, Impact Control proved extremely effective [60].

7.6 The Impact Transient

The transition from free space motion to contact with the environment provides the greatest test to the stability of the chosen control strategy. This is because of the almost instantaneous exertion of reaction forces upon the arm. Some researchers have addressed this problem by utilizing soft force sensors, or a soft ‘skin’ over the force sensor surface [72, 3]. The introduction of extra compliance extends the period of impact and absorbs some of the energy. As mentioned in Section 7.4, we have chosen not to use passive compliance because of its intrinsic problems such as mechanical fatigue and nonlinearities.

Considering the case of hard surface to hard surface impact, the transient time is very short. For a manipulator end-effector moving at 1 m/s impacting a surface with stiffness of 10^4 N/m, the force will initially increase at a rate of 10^4 N/s. If a resolution of 1 N is considered adequate for control, the sampling rate must be 10^4 Hz. This is 25 times faster than our sampling rate. Further, a required sampling rate of over 10^6 Hz would be necessary for robustness to some of the surfaces/speeds that were tried in our experimentation.

This impact transient is further complicated by the available torque of the actuators. Even if the sampling rate were fast enough to adequately detect the rise in external forces due to the impact, the joint torque necessary to substantially soften the impact is not available. In other words, the arm cannot stop itself instantaneously to prevent the impact. Assuming a maximum allowed impact force of 10 N, the arm would have to stop in 1 mm. To stop this suddenly, an acceleration of 500 m/s^2 is required (over 50 g!). For a manipulator with an effective Cartesian

Space mass of 1 kg, 500 N is required. At least 98% of this force must be provided by the arm itself. Obviously this is not feasible for conventional actuators and manipulators.

A third problem is the kinematic fluctuations of the arm during impact. This is mainly due to the compression and flexion caused by the impact forces. For instance, vibrations in the links will cause changes in the end effector position, although no change in joint position is measured. This can cause problems for schemes that rely on accurate measurement of the surface position or velocity, such as stiff impedance controllers [49] or impact damping strategies [28].

8 Conclusions

This chapter has presented solutions which address the three main issues of both space and terrestrial manipulation: collision free motion through the environment, stable transition through the contact phase, and accurate force control on the surface. To move through the environment successfully, an artificial force function based on superquadrics has been reviewed. This function can provide collision avoidance, or surface approach at safe velocities. After impact has occurred, stability may be maintained with the impact controller proposed. Subsequent force controlled manipulation is best achieved with integral gain force control. Also, though analysis of the impact and force control strategies, the essential equivalence of second order impedance control and proportional gain force control has been demonstrated. Finally, an in-depth discussion of the implementational consideration for these schemes has been provided.

9 Acknowledgments

This research was performed at Carnegie Mellon University and supported by an Air Force Graduate Laboratory Fellowship (for Richard Volpe), DARPA under contract DAAA-21-89C-0001, the Department of Electrical and Computer Engineering, and The Robotics Institute.

The writing and publication of this paper was supported by the above and the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

The views and conclusion contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Air Force, DARPA, or the U.S. Government. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

References

- [1] P. Akella, R. Siegwart, and M. Cutkosky. Manipulation with Soft Fingers: Contact Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 652–657, 1991.
- [2] B. Allotta et al. Controlling contact by integrating proximity and force sensing. In *Second International Conference on Experimental Robotics*, Toulouse, France, June 25-27 1991.
- [3] C. An and J. Hollerbach. Dynamic Stability Issues in Force Control of Manipulators. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 890–896, 1987.
- [4] J. R. Andrews and N. Hogan. Impedance control as a framework for implementing obstacle avoidance in a manipulator. In David E. Hardt and Wayne J. Book, editors, *Control Of Manufacturing Processes And Robotic Systems*, pages 243–251. ASME, 1983.
- [5] R. Bajcsy and F. Solina. Three dimensional object representation revisited. In *First International Conference on Computer Vision*, pages 231–240, London, England, June 8-11 1987.
- [6] A.H. Barr. Superquadrics and angle-preserving transformations. *IEEE Computer Graphics and Applications*, 1:11–23, 1981.
- [7] C. Boddy and J. Taylor. Whole-arm reactive collision avoidance control of kinematically redundant manipulators. In *IEEE International Conference on Robotics and Automation*, pages 382–387, Atlanta, Georgia, May 2-6 1992.
- [8] E. Cheung and V. Lumelsky. Motion planning for robot arm manipulators with proximity sensing. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 740–745, Philadelphia, Pennsylvania, April 24–29 1988.
- [9] B. Cohen-Tannoudji, C. Diu, and F. Laloe. *Quantum Mechanics*, volume 2. John Wiley and Sons, New York, 1977.
- [10] E. Colgate and N. Hogan. An Analysis of Contact Instability In Terms of Passive Physical Equivalents. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 404–409, 1989.
- [11] J. E. Colgate and N. Hogan. Robust control of dynamically interacting systems. *International Journal of Control*, 48(1):65–88, 1988.
- [12] J. De Schutter. A Study of Active Compliant Motion Control Methods For Rigid Manipulators Based on a Generic Scheme. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1060–1065, 1987.
- [13] J. De Schutter. Improved Force Control Laws For Advanced Tracking Applications. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1497–1502, 1988.
- [14] S. Eppinger and W. Seering. On Dynamic Models of Robot Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 29–34, 1986.
- [15] S. Eppinger and W. Seering. Understanding Bandwidth Limitations on Robot Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 904–909, Raleigh, N.C., 1987.

- [16] B. Faverjon. Obstacle avoidance using an octree in the configuration space of a manipulator. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 504–512, Atlanta, GA, 1984.
- [17] K. Fu, R. Gonzalez, and C. Lee. *Robotics: Control, Sensing, Vision, and Intelligence*. McGraw-Hill, New York, 1987.
- [18] A. Goldenberg. Implementation of Force and Impedance Control in Robot Manipulators. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1626–1632, 1988.
- [19] N. Hogan. Impedance control: An approach to manipulation. *Journal of Dynamic Systems, Measurement, and Control*, 107:1–24, March 1985.
- [20] N. Hogan. Impedance Control: An Approach to Manipulation: Parts I, II, and III. *Journal of Dynamic Systems, Measurement, and Control*, 107:1–24, March 1985.
- [21] N. Hogan. Stable Execution of Contact Tasks Using Impedance Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1047–1054, 1987.
- [22] N. Hogan and S.L. Cotter. Cartesian Impedance Control of a Nonlinear Manipulator. In W.J. Book, editor, *Robotics Research and Advanced Applications*, pages 121–128, New York, 1982. ASME.
- [23] H. Ishikawa, C. Sawada, K. Kawase, and M. Takata. Stable Compliance Control and Its Implementation for a 6 D.O.F. Manipulator. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 98–103, 1989.
- [24] J. Kahng and F. Amirouche. Impact Force Analysis in Mechanical Hand Design — Part I. *International Journal of Robotics and Automation*, 3(3):158–164, 1988.
- [25] H. Kazerooni. Robust, Non-Linear Impedance Control for Robot Manipulators. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 741–750, 1987.
- [26] H. Kazerooni, T. Sheridan, and P. Houpt. Robust Compliant Motion for Manipulators, Parts I and II. *IEEE Journal of Robotics and Automation*, RA-2(2):83–105, June 1986.
- [27] O. Khatib. Real-Time Obstacle Avoidance for Manipulators and Mobile Robots. *The International Journal of Robotics Research*, 5(1), 1986.
- [28] O. Khatib and J. Burdick. Motion and Force Control of Robot Manipulators. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1381–1386, 1986.
- [29] P. Khosla. *Real-Time Control and Identification of Direct Drive Manipulators*. PhD thesis, Carnegie Mellon University, Department of Computer and Electrical Engineering, August 1986.
- [30] Daniel E. Koditschek. Exact robot navigation by means of potential functions: Some topological considerations. In *IEEE International Conference on Robotics and Automation*, Raleigh, North Carolina, March 31 – April 3 1987.
- [31] B. Krogh. A generalized potential field approach to obstacle avoidance control. In *SME Conf. Proc. Robotics Research: The Next Five Years and Beyond*, Bethlehem, Pennsylvania, August 1984.

- [32] D. Lawrence. Impedance Control Stability Properties in Common Implementations. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1185–1190, 1988.
- [33] T. Lozano-Perez. Spatial planning: A configuration space approach. *IEEE Transactions on Computers*, C-32(2):102–120, 1983.
- [34] T. Lozano-Perez. A simple motion planning algorithm for general manipulators. Ai lab memo, MIT, AI Labs, Cambridge, MA, 1986.
- [35] T. Lozano-Perez and M. Wesley. An algorithm for planning collision free paths among polyhedral objects. *Comm. of ACM*, 22(10):560–570, 1979.
- [36] J. Maples and J. Becker. Experiments in Force Control of Robotic Manipulators. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 695–702, 1986.
- [37] F. Miyazaki and S. Arimoto. Sensory Feedback for Robot Manipulators. *Journal of Robotic Systems*, 2(1):53–71, 1985.
- [38] W. Newman. Automatic obstacle avoidance at high speeds via reflex control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1104–1109, 1989.
- [39] W. S. Newman and N. Hogan. High speed control and obstacle avoidance using dynamic potential functions. In *IEEE International Conference on Robotics and Automation*, Raleigh, North Carolina, March 31 – April 3 1987.
- [40] C. Nguyen, S. Antrazi, and C. Campbell. Virtual Force Control of a SixDegreeofFreedom Parallel Manipulator. In *Robotics and Manufacturing: Recent Trends in Research, Education, and Applications*, pages 357–363. ASME Press, 1992.
- [41] J. Novak and J. Feddema. A capacitance-based proximity sensor for whole arm obstacle avoidance. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1307–1314, 1992.
- [42] M. Okutomi and M. Mori. Decision of robot movement by means of a potential field. *Advanced Robotics*, 1(2):131–141, 1986.
- [43] E. Papadopoulos and S. Dubowsky. On the Nature of Control Algorithms for Free-Floating Space Manipulators. *IEEE Transactions on Robotics and Automation*, 7(6):750–758, 1991.
- [44] R. Paul. Problems and Research Issues Associated with the Hybrid Control of Force and Displacement. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1966–1971, June 1987.
- [45] R. Paul and C. Wu. Manipulator Compliance Based on Joint Torque. In *IEEE Conference on Decision and Control*, pages 88–94, New Mexico, 1980.
- [46] M. Raibert and J. Craig. Hybrid Position/Force Control of Manipulators. *Journal of Dynamic Systems, Measurement, and Control*, 103(2):126–133, June 1981.
- [47] E. Rimon and E. Koditschek. The construction of analytic diffeomorphisms for exact robot navigation on star worlds. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 21–26, 1989.

- [48] R. K. Roberts. *The Compliance of End Effector Force Sensors for Robot Manipulator Control*. PhD thesis, Purdue University, Department of Mechanical Engineering, December 1984.
- [49] J. K. Salisbury. Active Stiffness Control of a Manipulator in Cartesian Coordinates. In *IEEE Conference on Decision and Control*, pages 95–100, New Mexico, 1980.
- [50] J. Schwartz and M. Sharir. On the piano movers problem, part ii. Technical Report 41, Courant Institute of Mathematical Sciences, NY, NY, 1982.
- [51] A. Sharon, N. Hogan, and D. Hardt. Controller Design in the Physical Domain (Application to Robot Impedance Control). In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 552–559, 1989.
- [52] W. Townsend and J. Salisbury. The Effect of Coulomb Friction and Stiction on Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 883–889, 1987.
- [53] S. Udupa. Collision detection and avoidance in computer controlled manipulators. In *Proceedings of the 5-th Joint International Conference on AI*. AI, 1977.
- [54] Z. Vafa and S. Dubowsky. On the Dynamics of Space Manipulators Using the Virtual Manipulator, with Applications to Path Planning. *The Journal of the Astronautical Sciences*, 38(4):441–472, 1990.
- [55] D. Vischer and O. Khatib. Design and Development of Torque-Controlled Joints. In V. Hayward and O. Khatib, editors, *Experimental Robotics I*, pages 271–286. Springer-Verlag Berlin Heidelberg, 1990.
- [56] R. Volpe. *Real and Artificial Forces in the Control of Manipulators: Theory and Experiments*. PhD thesis, Carnegie Mellon University, Department of Physics, September 1990.
- [57] R. Volpe and P. Khosla. A Strategy For Obstacle Avoidance and Approach Using Superquadric Potential Functions. In *Proceedings of the 5th International Symposium of Robotics Research*, pages 93–100, Tokyo, Japan, August 28-31 1989. MIT Press.
- [58] R. Volpe and P. Khosla. Manipulator Control with Superquadric Artificial Potential Functions: Theory and Experiments. *IEEE Transactions on Systems, Man, and Cybernetics; Special Issue on Unmanned Vehicles and Intelligent Systems*, November/December 1990.
- [59] R. Volpe and P. Khosla. Theoretical Analysis and Experimental Verification of a Manipulator / Sensor / Environment Model for Force Control. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, Los Angeles, November 1990.
- [60] R. Volpe and P. Khosla. Experimental Verification of a Strategy for Impact Control. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Sacramento, CA, April 1991.
- [61] R. Volpe and P. Khosla. The Equivalence of Second Order Impedance Control and Proportional Gain Explicit Force Control: Theory and Experiments. In *Proceedings of the Second Annual International Symposium on Experimental Robotics*, Toulouse, France, June 1991.
- [62] R. Volpe and P. Khosla. An Analysis of Manipulator Force Control Strategies Applied to an Experimentally Derived Model. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Raleigh, North Carolina, July 7-10 1992.

- [63] R. Volpe and P. Khosla. An Experimental Evaluation and Comparison of Explicit Force Control Strategies for Robotic Manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Nice, France, May 10-15 1992.
- [64] R. Volpe and P. Khosla. A Theoretical and Experimental Investigation of Explicit Force Control Strategies for Manipulators. *IEEE Transactions on Automatic Control*, 38(11), November 1993.
- [65] R. Volpe and P. Khosla. A Theoretical and Experimental Investigation of Impact Control for Manipulators. *International Journal of Robotics Research*, 12(4):351–365, August 1993.
- [66] J. Vranish and D. Chauhan. Tri-Mode Collision Avoidance Skin for Robot Arms in Space. In *Robotics and Manufacturing: Recent Trends in Research, Education, and Applications*, pages 189–195, British Columbia, Canada, August 28-31 1990. ASME Press.
- [67] C. Warren. Global path planning using artificial potential fields. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 316–321, 1989.
- [68] D. Wedel and Saridis G. An Experiment in Hybrid Position/Force Control of a Six DOF Revolute Manipulator. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1638–1642, 1988.
- [69] D. Whitney. Force Feedback Control of Manipulator Fine Motions. *Journal of Dynamic Systems, Measurement, and Control*, pages 91–97, June 1977.
- [70] D. Whitney. Historical Perspective and State of the Art in Robot Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 262–268, 1985.
- [71] Y. Xu and T. (editors) Kande. *Space Robotics: Dynamics and Control*. Kluwer Academic Publishers, 1993.
- [72] Y. Xu and R. Paul. On Position Compensation and Force Control Stability of a Robot with a Compliant Wrist. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 1173–1178, 1988.
- [73] K. Yamada. Arm Path Planning for a Space Robot. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, pages 2049–2055, July 26-30 1993.
- [74] K. Yoshida and N. Sashida. Modeling Impact Dynamics and Impulse Minimization for Space Robots. In *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems*, pages 2064–2069, July 26-30 1993.
- [75] K. Youcef-Toumi. Force Control of Direct-Drive Manipulators For Surface Following. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 2055–2060, 1987.
- [76] K. Youcef-Toumi and D. Gutz. Impact and Force Control. In *Proceedings of the IEEE Conference on Robotics and Automation*, pages 410–416, 1989.